

Kernel methods are a simple and efficient way to extend linear statistical models to more flexible nonlinear models. They are especially useful in problems where the data is high-dimensional and heterogeneous, which is often the case when dealing with medical images, genetics, clinical data, etc. In this tutorial, we will cover the basics of how kernel methods work and see how they can be applied to discrimination and regression problems in medical imaging studies. The core idea behind kernel methods is to nonlinearly embed data into a derived feature space and then apply standard linear methods. The so-called "kernel trick" is that this embedding does not have to be computed explicitly. Rather, we only need to know how to take inner products between pairs of data points.